

An N -Body Solution to the Problem of Fock Exchange

Matt Challacombe* and Nicolas Bock

Theoretical Division, Los Alamos National Laboratory

We report an N -Body approach to computing the Fock exchange matrix with and without permutational symmetry. The method achieves an $O(N \lg N)$ computational complexity through an embedded metric-query, allowing hierarchical application of direct SCF criteria. The advantages of permutational symmetry are found to be 4-fold for small systems, but decreasing with increasing system size and/or more permissive neglect criteria. This work sets the stage for: (1) the introduction of range queries in multi-level multipole schemes for rank reduction, and (2) recursive task parallelism.

INTRODUCTION

For physical problems that find compact representation supported by fast transforms, such as the fast wavelet transform and the fast Fourier transform, recursion and reduced complexity are intrinsic. For problems that do not find compact representation, there may be fast N -Body solutions. N -Body solvers combine recursive subdivision, elements of database theory that include efficient metric- and range-queries, as well as multi-level approaches to rank reduction and culling. In addition to a reduced computational complexity, these database elements enable the exploitation of temporal and spatial locality in high performance implementations, making the N -Body programming model one of the most successful in large scale scientific simulation [1–9]. So far, it has been possible to develop reduced complexity, N -Body solutions for many aspects of conventional quantum chemical self consistent field theory [10], including the Hartree problem [11–14], exchange correlation cubature [15], inverse factorization and spectral projection based on the SpAMM algorithm [16, 17].

Beyond the recursive multi-wavelet work of Yanai *et al.* [18], N -Body solutions to Fock exchange in the conventional, Gaussian Type Atomic Orbital (GTAO) representation remains an open and important problem. It's important because the Fock exchange is a key ingredient in hybrid theories [19, 20] that yield qualitatively better results for many challenging problems (relative to “pure” DFT), from metal oxides [21–28], to chemical reactions [29, 30] battery materials [31], photovoltaic semiconductors [32], and in biochemistry [33–35].

Reduced complexity approaches to GTAO Fock exchange, such as ONX and its variants [36–43], remain predicated on density matrix truncation (sparsification), preordered skip-out lists and/or matrix range estimates, as well as list-of-lists matrix structures such as BCSR [44]. These preprocessing steps and their associated data structures greatly complicate aspects of the problem involving range-queries and metric-queries necessary to implement multi-level methods for rank reduction and culling [41, 45], and also force a choice between “integral

driven” and “index driven” schemes for domain decomposition [46, 47]. In this contribution, we outline an N -body reformulation of Fock exchange matrix construction that enables task decomposition in the recursive space of tensor products, and which also provides an embedded framework for metric- and range-queries.

RECURSIVE QUANTUM CHEMISTRY

We begin our development with recursive bisection of the indicial space of GTAO basis functions:

$$|\mu\rangle^k = \frac{|\mu_0\rangle^{k+1}}{|\mu_1\rangle^{k+1}} = \frac{\begin{bmatrix} \phi_{\mu_0^L} \\ \phi_{\mu_0^{L+1}} \\ \vdots \\ \phi_{\mu_0^R} \\ \phi_{\mu_1^L} \\ \phi_{\mu_1^{L+1}} \\ \vdots \\ \phi_{\mu_1^R} \end{bmatrix}}{1}, \quad (1)$$

where $|\mu\rangle^k$ denotes a vector (block) of functions ϕ at depth k with span $\mu \in [\mu^L, \mu^R]$. This bisection can be accomplished in a variety of ways. Here, we consider bisection with ragged edges, rather than by powers of two, so that the underlying atomic-orbital shell structure (*ie.* $p, sp, d, f \dots$) is preserved, greatly simplifying the associated computation of two-electron integrals.

The fundamental telescoping quantities are matrix quadrees, discussed in Refs. [16, 17] and references therein, and shell pairs, quadrees that obtain recursively from shell-shell outerproducts:

$$\begin{aligned} |\mu\nu\rangle^k &= \begin{pmatrix} |\mu_0\nu_0\rangle^{k+1} & |\mu_0\nu_1\rangle^{k+1} \\ |\mu_1\nu_0\rangle^{k+1} & |\mu_1\nu_1\rangle^{k+1} \end{pmatrix} \\ &= \begin{pmatrix} |\mu_0\rangle^{k+1} \otimes |\nu_0\rangle^{k+1} & |\mu_0\rangle^{k+1} \otimes |\nu_1\rangle^{k+1} \\ |\mu_1\rangle^{k+1} \otimes |\nu_0\rangle^{k+1} & |\mu_1\rangle^{k+1} \otimes |\nu_1\rangle^{k+1} \end{pmatrix} \end{aligned} \quad (2)$$

In the large system limit, the complexity with respect to number of GTAO basis functions N becomes $\mathcal{O}(N)$ due to non-overlapping functions,

$$|\mu\nu\rangle^k = \begin{cases} 0 & \text{if } |\mu\rangle^k \cap |\nu\rangle^k = \emptyset \\ |\mu\rangle^k \otimes |\nu\rangle^k & \text{else} \end{cases} \quad (3)$$

where non-intersection obtains if all overlap integrals between shell-blocks are sufficiently small,

$$|\mu\rangle^k \cap |\nu\rangle^k = \emptyset \quad \text{if} \quad (\phi_\mu, \phi_\nu) < \tau_{\text{ovlp}} \\ \forall \mu \in [\mu^L, \mu^R], \nu \in [\nu^L, \nu^R], \quad (4)$$

as determined by an overlap threshold τ_{ovlp} .

EXCHANGE AS HEXTREE TRAVERSAL

At the top level ($k = 0$), the Fock exchange matrix can be written simply with the less than orthodox bra-ket notation:

$$\mathbf{K}^0 = -\frac{1}{2} (\mu_0 \nu_0 |^0 \mathbf{P}^0 | \lambda_0 \sigma_0)^0, \quad (5)$$

which is useful shorthand for the tensor contraction

$$K_{\mu\sigma} = -\frac{1}{2} \sum_{\nu,\lambda} P_{\nu,\lambda} (\mu\nu | \lambda\sigma), \quad (6)$$

where $(\mu\nu | \lambda\sigma)$ is the standard two-electron integral over GTAO basis functions [10].

In the case of naive recursion, where permutational symmetry of the two-electron integrals is unexploited, shell-pair quadrees maintain their relationship with sub-indices on recursion, permitting the simplified notation: $|\mathbf{00}\rangle^k \equiv |\mu_0 \nu_0\rangle^k$, $|\mathbf{01}\rangle^k \equiv |\mu_0 \nu_1\rangle^k$ and so on. Then, at all levels, sub-blocks of the Fock exchange matrix are

$$\mathbf{K}_{\mathbf{00}}^k \leftarrow \mathbf{K}_{\mathbf{00}}^k - \left[(\mathbf{00} |^k \mathbf{P}_{\mathbf{00}}^k | \mathbf{00})^k + (\mathbf{00} |^k \mathbf{P}_{\mathbf{01}}^k | \mathbf{10})^k \right. \\ \left. + (\mathbf{10} |^k \mathbf{P}_{\mathbf{10}}^k | \mathbf{00})^k + (\mathbf{01} |^k \mathbf{P}_{\mathbf{11}}^k | \mathbf{10})^k \right] / 2 \quad (7a)$$

$$\mathbf{K}_{\mathbf{01}}^k \leftarrow \mathbf{K}_{\mathbf{01}}^k - \left[(\mathbf{00} |^k \mathbf{P}_{\mathbf{00}}^k | \mathbf{01})^k + (\mathbf{01} |^k \mathbf{P}_{\mathbf{10}}^k | \mathbf{01})^k \right. \\ \left. + (\mathbf{00} |^k \mathbf{P}_{\mathbf{01}}^k | \mathbf{11})^k + (\mathbf{01} |^k \mathbf{P}_{\mathbf{11}}^k | \mathbf{11})^k \right] / 2 \quad (7b)$$

$$\mathbf{K}_{\mathbf{10}}^k \leftarrow \mathbf{K}_{\mathbf{10}}^k - \left[(\mathbf{10} |^k \mathbf{P}_{\mathbf{00}}^k | \mathbf{00})^k + (\mathbf{11} |^k \mathbf{P}_{\mathbf{10}}^k | \mathbf{00})^k \right. \\ \left. + (\mathbf{10} |^k \mathbf{P}_{\mathbf{01}}^k | \mathbf{10})^k + (\mathbf{11} |^k \mathbf{P}_{\mathbf{11}}^k | \mathbf{10})^k \right] / 2 \quad (7c)$$

$$\mathbf{K}_{\mathbf{11}}^k \leftarrow \mathbf{K}_{\mathbf{11}}^k - \left[(\mathbf{10} |^k \mathbf{P}_{\mathbf{00}}^k | \mathbf{01})^k + (\mathbf{11} |^k \mathbf{P}_{\mathbf{10}}^k | \mathbf{01})^k \right. \\ \left. + (\mathbf{10} |^k \mathbf{P}_{\mathbf{01}}^k | \mathbf{11})^k + (\mathbf{11} |^k \mathbf{P}_{\mathbf{11}}^k | \mathbf{11})^k \right] / 2 \quad (7d)$$

equivalent to hextree traversal in the recursive task space of exchange tensor contraction. An advantage of this construction is that, together with quadrees that are

informed at each level by the Frobenius norm $\|\cdot\|_F$, the blocked Almlöf-Alhrichs criteria [48, 49]

$$\|(\mu\nu | \mu\nu)^k\|_F \cdot \|\mathbf{P}_{\nu\lambda}^k\|_F \cdot \|(\lambda\sigma | \lambda\sigma)^k\|_F \leq \tau_{2e}, \quad (8)$$

can be carried out naturally via embedded metric-query, enabling recursion termination when the bound is satisfied. Because $\|\cdot\|_F$ is sub-multiplicative, this procedure is rigorously equivalent to the standard direct SCF method, with the integral threshold τ_{2e} retaining its conventional meaning.

The efficiency of this query in culling negligible integral contributions is dependent on the numerical structure of the underlying data. A simple solution to this problem involves ordering shells with a locality preserving space filling curve, effectively clustering elements of like magnitude as shown in Figs. 1 and 2 of Ref. [17]. Note that methods relying on random permutation to make use of Cannon's algorithm for the parallel multiplication of sparse matrices [50, 51] destroy these locality properties. As with the SpAMM [16, 17] solver, truncation of the vector space (sparsification) is not a prerequisite for achieving reduced complexities, but well structured density matrices with decay are.

RECURSION WITH SYMMETRY

Extending the recursive approach outlined in the previous section to the exploitation of 4-fold permutational symmetry is more involved than in conventional schemes [39, 43]. Naively, we wish to employ

$$\mathbf{K}_{\mu\sigma}^k \leftarrow \mathbf{K}_{\mu\sigma}^k - (\mu\nu |^k \mathbf{P}_{\nu\lambda}^k | \lambda\sigma)^k / 2 \quad (9a)$$

$$\mathbf{K}_{\nu\sigma}^k \leftarrow \mathbf{K}_{\nu\sigma}^k - (\nu\mu |^k \mathbf{P}_{\mu\lambda}^k | \lambda\sigma)^k / 2 \quad (9b)$$

$$\mathbf{K}_{\mu\lambda}^k \leftarrow \mathbf{K}_{\mu\lambda}^k - (\mu\nu |^k \mathbf{P}_{\nu\sigma}^k | \sigma\lambda)^k / 2 \quad (9c)$$

$$\mathbf{K}_{\nu\lambda}^k \leftarrow \mathbf{K}_{\nu\lambda}^k - (\nu\mu |^k \mathbf{P}_{\mu\sigma}^k | \sigma\lambda)^k / 2 \quad (9d)$$

which has the potential to yield speedups of up to 4x in the evaluation of two-electron integrals, but which does not reduce the cost of tensor contraction. To avoid over-computing however, the blockwise restriction

$$\mu \leq \nu, \quad \lambda \leq \sigma, \quad \mu + \nu(\nu - 1)/2 \leq \lambda + \sigma(\sigma - 1)/2 \\ \forall \mu \in [\mu^L, \mu^R], \quad \nu \in [\nu^L, \nu^R], \quad (10) \\ \lambda \in [\lambda^L, \lambda^R], \quad \text{and} \quad \sigma \in [\sigma^L, \sigma^R].$$

must be observed at each level of recursion. Satisfying these inequalities without carrying along an explosion of auxiliary source and sink sub-matrices requires that we untether the strict 1-to-1 relationship between target symmetries and matrix sub-indices, putting the transpose operation into play for shell-pairs and matrices. These complications are resolved by introducing an intermediate level of recursion, where target symmetries




Figure 1: Onset of linear scaling for the symmetry enhanced N -Body Fock exchange solver using two screening regimes: $(\tau_{2e}, \tau_{ovlp}) = (1d-8, 1d-11)$ and $(1d-10, 1d-13)$. The inset is the ratio of CPU times for naive recursion relative to symmetry enhanced recursion.

are determined and links to density and exchange matrix sub-blocks are set, up to 4 of 16 possible links each. Including Eq. (9), there are 8 conditions that exploit the full 4-fold symmetry of Fock exchange under recursion (the remaining 7 are given in the Appendix), which involve non-standard, “across the bar” (transpose) permutations as well as additional factors of two, (*e.g.* due to restricting the density matrix above the diagonal). In addition to cases of 4-fold symmetry, there are many other intermediate and terminal cases where there are fewer than 4 sub-matrices involved in recursion, for example due to cases where blockwise symmetry operations do not yield a full 4-fold complement, and also in cases where the density matrix is sparse and sub-blocks are not available. These later cases that fall outside of Eqs. (9) & (A.11-A.17) are referred to as “sparse” in the following. Also, the blockwise culling of negligible integral contributions is carried out recursively, as in Eq. (8), but instead using the maximum density matrix norms as they occur at each level.

IMPLEMENTATION AND RESULTS

Recursive construction of the Fock exchange matrix was implemented with and without permutational symmetry in a development version of **FreeON** [52], using the Head-Gordon Pople algorithm for two-electron integrals [53]. Non-sparse contractions, involving full 4-fold matrix compliments, are carried out in a single code block, different from conventional contraction schemes in that there is no auxiliary “batch” dimension available for vector optimization. In the current implementation, recur-




Figure 2: Percentage breakdown by contraction blocks, Eqs. (9) & (A.11-A.17), including also “sparse” blocks, in the symmetry enhanced N -Body Fock exchange solver in the tighter $(1d-10, 1d-13)$ regime. Inset shows blocks with fewer occurrences, corresponding to Eqs. (A.12-A.17).

sion extends to blocks with dimension 10×10 or smaller. All calculations were carried out on an Intel Xeon CPU E5-2687W @ 3.10GHz using v. 14.0.1 of the Intel Fortran and C compilers.

Results are shown in Fig. 1 and 2 for the standard sequence of water droplets corresponding to STP conditions, $(H_2O)_n$ with $n = 10, 30, 50, 70, 90, 110, 130$, and density matrices tightly converged using an atom blocked threshold $\tau_{mtrx} = 1d-6$ [44] at the B3LYP/6-31G** level of theory. All systems were ordered using the Hilbert-curve to maximize locality as shown in Figs. 1 and 2 of Ref. [17]. In the largest droplet, $(H_2O)_{130}$, ragged bisection with $\sim 10 \times 10$ blocks yields a depth $k = 10$, comprising $\sim 6d+5$ nodes in the shell-pair quadtree.

In Fig. 1, an early approach to linear scaling is shown for two values of the integral screening parameter τ_{2e} , together with the relative performance of the naive vs symmetry enhanced methods as inset. In the case of permutational symmetry, it is found that the cost of integral evaluation and contraction tend towards 1:1. Figure 2 shows the % occurrence by symmetry block, Eqs. (9) & (A.11-A.17), for the symmetry enhanced method. The dominant occurrences correspond to Eq. (9) and (A.11) at about 60% and 20% respectively, followed by the sparse case at 10%. Because this recursive task space corresponds to a hextree, there are ample opportunities for Certainly, the task parallel features of **OpenMP 3.0** [54] are ideally suited to this formulation, and middleware for distributed memory task parallelism, such as **charm++** [55], are becoming increasingly powerful and simple to use. Also, the cost of integral contraction with and without symmetry is about 2:3, reflecting loop overheads and higher levels of optimization that favor the combined

4-fold contraction blocks.

SUMMARY

We’ve presented a novel, N -Body formulation for the naive and symmetry enhanced construction of the Fock exchange matrix, which achieves a reduced, linear scaling complexity for density matrices with decay. The method does not employ any of the matrix truncation (sparsification), integral skipout lists, matrix range estimation or list-of-lists data structures employed by conventional methods. Also, the method may enable transcendence of conventional “index driven” vs “integral” driven paradigms for parallel Fock matrix construction [46, 47], through decomposition in the recursive task space of the tensor contraction. Even though this task space is sparse and irregular due to culling, because it is higher dimensional, corresponding to hextree traversal, it offers ample opportunity for parallelism through well developed middleware such as `OpenMP 3.0` [54] in shared memory environments and `charm++` [55] in distributed environments.

With this contribution, all minimally essential components of reduced complexity electronic structure theory at the self consistent field level of theory have been reformulated as N -Body solvers. It remains to be seen how tightly these solvers can be integrated using a common infrastructure and programming model. Finally, it also remains to be seen if the ability to embed range- and metric-queries into this structure can be exploited to achieve true multi-level rank reduction for computation of the Fock exchange matrix.

This work was supported by the U. S. Department of Energy under Contract No. DE-AC52-06NA25396 and LDRD-ER grant 20110230ER. The authors acknowledge the stimulating environment of the Ten-Bar Café and expert administrations of the barista.

* Electronic address: matt.challacombe@freeon.org; corresponding author; URL: <http://www.freeon.org>

- [1] M. Warren and J. Salmon, in *Proc. Supercomput. '92* (1992), pp. 570–576.
- [2] M. Taiji, J. Makino, T. Ebisuzaki, and D. Sugimoto, in *Proc. 8th Int. Parallel Process. Symp.* (1994).
- [3] T. Fukushige and J. Makino, in *Proc. 1996 ACM/IEEE Conf. Supercomput. - Supercomput. '96* (ACM Press, New York, New York, USA, 1996), pp. 48–es.
- [4] M. Warren, J. Salmon, D. Becker, M. Goda, T. Sterling, and W. Winckelmans, *ACM/IEEE SC 1997 Conf.* (1997).
- [5] A. Kawai, T. Fukushige, and J. Makino, in *Proc. 1999 ACM/IEEE Conf. Supercomput. - Supercomput. '99* (ACM Press, New York, New York, USA, 1999), pp. 67–es.
- [6] J. Makino, E. Kokubo, and T. Fukushige, in *Proc. 2003 ACM/IEEE Conf. Supercomput. - SC '03* (ACM Press, New York, New York, USA, 2003), p. 2.
- [7] T. Hamada, T. Narumi, R. Yokota, K. Yasuoka, K. Nitadori, and M. Taiji, in *Proc. Conf. High Perform. Comput. Networking, Storage Anal. - SC '09* (ACM Press, New York, New York, USA, 2009), p. 1.
- [8] T. Hamada and K. Nitadori, in *2010 ACM/IEEE Int. Conf. High Perform. Comput. Networking, Storage Anal.* (IEEE, 2010), pp. 1–9.
- [9] T. Ishiyama, K. Nitadori, and J. Makino, in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (IEEE Computer Society Press, Los Alamitos, CA, USA, 2012), SC '12, pp. 5:1–5:10.
- [10] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry* (McGraw Hill, New York, 1982).
- [11] C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, *Chem. Phys. Lett.* **253**, 268 (1996).
- [12] M. Challacombe, E. Schwegler, and J. Almlöf, *The Journal of Chemical Physics* **104**, 4685 (1996).
- [13] M. Challacombe and E. Schwegler, *J. Chem. Phys.* **106**, 5526 (1997).
- [14] M. Challacombe, E. Schwegler, and J. Almlöf, in *Computational Chemistry: Reviews of Current Trends*, edited by J. Leszczynski (World Scientific, Singapore, 1996), vol. 1 of *Computational Chemistry: Reviews of Current Trends*, pp. 53–107.
- [15] M. Challacombe, *J. Chem. Phys.* **113**, 10037 (2000).
- [16] M. Challacombe and N. Bock, *arXiv Prepr. arXiv1011.3534* (2010).
- [17] N. Bock and M. Challacombe, *SIAM J. Sci. Comput.* **35**, C72 (2013).
- [18] T. Yanai, G. I. Fann, Z. Gan, R. J. Harrison, and G. Beylkin, *J. Chem. Phys.* **121**, 6680 (2004).
- [19] A. D. Becke, *The Journal of Chemical Physics* **98**, 1372 (1993).
- [20] A. D. Becke, *The Journal of Chemical Physics* **98**, 5648 (1993).
- [21] J. P. Perdew and M. Levy, *Phys. Rev. Lett.* **51**, 1884 (1983).
- [22] L. J. Sham and M. Schlüter, *Phys. Rev. B* **32**, 3883 (1985).
- [23] H. Xiao, J. Tahir-Kheli, and W. A. Goddard, *The Journal of Physical Chemistry Letters* **2**, 212 (2011).
- [24] T. Bredow and A. R. Gerson, *Phys. Rev. B* **61**, 5194 (2000).
- [25] B. Meyer and D. Marx, *Phys. Rev. B* **67**, 035403 (2003).
- [26] J. Muscat, A. Wander, and N. Harrison, *Chemical Physics Letters* **342**, 397 (2001).
- [27] F. Wang, C. Di Valentin, and G. Pacchioni, *The Journal of Physical Chemistry C* (2011).
- [28] N. Wilson, Ph.D. thesis, School of Applied Sciences, RMIT University (2009).
- [29] D. R. B. Brittain, C. Y. Lin, A. T. B. Gilbert, E. I. Izgorodina, P. M. W. Gill, and M. L. Coote, *Phys. Chem. Chem. Phys.* **11**, 1138 (2009).
- [30] Y. Zhao and D. G. Truhlar, *Journal of Chemical Theory and Computation* **7**, 669 (2011).
- [31] M. Ramzan, S. Lebeëgue, T. Kang, and R. Ahuja, *The Journal of Physical Chemistry C* (2011).
- [32] H. Xiao, J. Tahir-Kheli, and W. Goddard III, *The Journal of Physical Chemistry Letters* **2**, 212 (2011).
- [33] W. Ames and S. Larsen, *Journal of Biological Inorganic*

Chemistry **14**, 547 (2009), 10.1007/s00775-009-0469-9.

- [34] L. Rao, H. Ke, G. Fu, X. Xu, and Y. Yan, Journal of Chemical Theory and Computation **5**, 86 (2008).
- [35] M. Marianski, A. Asensio, and J. Dannenberg, The Journal of chemical physics **137**, 044109 (2012).
- [36] E. Schwegler and M. Challacombe, J. Chem. Phys. **105**, 2726 (1996).
- [37] E. Schwegler, M. Challacombe, and M. Head-Gordon, J. Chem. Phys. **106**, 9708 (1997).
- [38] E. Schwegler, M. Challacombe, and M. Head-Gordon, J. Chem. Phys. **109**, 8764 (1998).
- [39] C. Ochsenfeld, C. A. White, and M. Head-Gordon, J. Chem. Phys. **109**, 1663 (1998).
- [40] E. Schwegler and M. Challacombe, J. Chem. Phys. **111**, 6223 (1999).
- [41] E. Schwegler and M. Challacombe, J. Chem. Phys. **111**, 6223 (1999).
- [42] E. Schwegler, Ph.D. thesis, University of Minnesota (1998).
- [43] E. Schwegler and M. Challacombe, Theor. Chem. Accounts Theory, Comput. Model. (Theoretica Chim. Acta) **104**, 344 (2000).
- [44] M. Challacombe, Comp. Phys. Comm. **128**, 93 (2000).
- [45] S. A. Maurer, D. S. Lambrecht, D. Flaig, and C. Ochsenfeld, J. Chem. Phys. **136**, 144107 (2012).
- [46] H. Takashima, S. Yamada, S. Obara, K. Kitamura, S. Inabata, N. Miyakawa, K. Tanabe, and U. Nagashima, J. Comput. Chem. **23**, 1337 (2002).
- [47] V. Weber and M. Challacombe, J. Chem. Phys. **125**, 104110 (2006).
- [48] J. Almlöf, K. Faegri, and K. Korsell, Journal of Computational Chemistry **3**, 385 (1982).
- [49] M. Häser and R. Ahlrichs, Journal of Computational Chemistry **10**, 104 (1989).
- [50] A. Buluç and J. R. Gilbert, in *ICPP '08: Proceedings of the 2008 37th International Conference on Parallel Processing* (IEEE Computer Society, Washington, DC, USA, 2008), pp. 503–510.
- [51] A. Buluc and J. Gilbert, Arxiv preprint arXiv:1109.3739 (2011).
- [52] N. Bock, M. Challacombe, C. K. Gan, G. Henkelman, K. Nemeth, A. M. N. Niklasson, A. Odell, E. Schwegler, C. J. Tymczak, and V. Weber, *FreeON: A suite of programs for linear scaling quantum chemistry, branch "beta" v. a01a4e8bd80b87501cd31404818d2445c5469e61* (2014), <http://www.freeon.org/>.
- [53] M. Head-Gordon and J. A. Pople, J. Chem. Phys. **89**, 5777 (1988).
- [54] A. Duran, X. Teruel, R. Ferrer, X. Martorell, and E. Ayguade, in *Parallel Processing, 2009. ICPP '09. International Conference on* (2009), pp. 124–131.
- [55] L. V. Kale and S. Krishnan, in *Proceedings of the Conference on Object Oriented Programming Systems, Languages and Applications* (1993), pp. 91–108.

Appendix

Additional symmetry blocks:

$$\mathbf{K}_{\mu\sigma}^k \leftarrow \mathbf{K}_{\mu\sigma}^k - (\mu\nu|^k \mathbf{P}_{\nu\lambda}^k |\lambda\sigma)^k / 2 \quad (\text{A.11a})$$

$$\mathbf{K}_{\nu\sigma}^k \leftarrow \mathbf{K}_{\nu\sigma}^k - (\nu\mu|^k \mathbf{P}_{\mu\lambda}^k |\lambda\sigma)^k / 2 \quad (\text{A.11b})$$

$$\mathbf{K}_{\lambda\mu}^k \leftarrow \mathbf{K}_{\lambda\mu}^k - (\lambda\sigma|^k \mathbf{P}_{\sigma\nu}^k |\nu\mu)^k / 2 \quad (\text{A.11c})$$

$$\mathbf{K}_{\nu\lambda}^k \leftarrow \mathbf{K}_{\nu\lambda}^k - (\nu\mu|^k \mathbf{P}_{\mu\sigma}^k |\sigma\lambda)^k / 2 \quad (\text{A.11d})$$

$$\mathbf{K}_{\mu\sigma}^k \leftarrow \mathbf{K}_{\mu\sigma}^k - (\mu\nu|^k \mathbf{P}_{\nu\lambda}^k |\lambda\sigma)^k / 2 \quad (\text{A.12a})$$

$$\mathbf{K}_{\nu\lambda}^k \leftarrow \mathbf{K}_{\nu\lambda}^k - (\nu\mu|^k \mathbf{P}_{\mu\sigma}^k |\sigma\lambda)^k / 2 \quad (\text{A.12b})$$

$$\mathbf{K}_{\mu\lambda}^k \leftarrow \mathbf{K}_{\mu\lambda}^k - (\mu\nu|^k \mathbf{P}_{\nu\sigma}^k |\sigma\lambda)^k / 2 \quad (\text{A.12c})$$

$$\mathbf{K}_{\lambda\nu}^k \leftarrow \mathbf{K}_{\lambda\nu}^k - (\lambda\sigma|^k \mathbf{P}_{\sigma\mu}^k |\mu\nu)^k / 2 \quad (\text{A.12d})$$

$$\mathbf{K}_{\mu\sigma}^k \leftarrow \mathbf{K}_{\mu\sigma}^k - (\mu\nu|^k \mathbf{P}_{\nu\lambda}^k |\lambda\sigma)^k / 2 \quad (\text{A.13a})$$

$$\mathbf{K}_{\nu\sigma}^k \leftarrow \mathbf{K}_{\nu\sigma}^k - (\nu\mu|^k \mathbf{P}_{\mu\lambda}^k |\lambda\sigma)^k / 2 \quad (\text{A.13b})$$

$$\mathbf{K}_{\mu\lambda}^k \leftarrow \mathbf{K}_{\mu\lambda}^k - (\mu\nu|^k \mathbf{P}_{\nu\sigma}^k |\sigma\lambda)^k / 2 \quad (\text{A.13c})$$

$$\mathbf{K}_{\lambda\nu}^k \leftarrow \mathbf{K}_{\lambda\nu}^k - (\lambda\sigma|^k \mathbf{P}_{\sigma\mu}^k |\mu\nu)^k / 2 \quad (\text{A.13d})$$

$$\mathbf{K}_{\mu\sigma}^k \leftarrow \mathbf{K}_{\mu\sigma}^k - (\mu\nu|^k \mathbf{P}_{\nu\lambda}^k |\lambda\sigma)^k / 2 \quad (\text{A.14a})$$

$$\mathbf{K}_{\nu\sigma}^k \leftarrow \mathbf{K}_{\nu\sigma}^k - (\nu\mu|^k \mathbf{P}_{\mu\lambda}^k |\lambda\sigma)^k / 2 \quad (\text{A.14b})$$

$$\mathbf{K}_{\lambda\mu}^k \leftarrow \mathbf{K}_{\lambda\mu}^k - (\lambda\sigma|^k \mathbf{P}_{\sigma\nu}^k |\nu\mu)^k / 2 \quad (\text{A.14c})$$

$$\mathbf{K}_{\lambda\nu}^k \leftarrow \mathbf{K}_{\lambda\nu}^k - (\lambda\sigma|^k \mathbf{P}_{\sigma\mu}^k |\mu\nu)^k / 2 \quad (\text{A.14d})$$

$$\mathbf{K}_{\mu\sigma}^k \leftarrow \mathbf{K}_{\mu\sigma}^k - (\mu\nu|^k \mathbf{P}_{\nu\nu}^k |\nu\sigma)^k / 2 \quad (\text{A.15a})$$

$$\mathbf{K}_{\nu\sigma}^k \leftarrow \mathbf{K}_{\nu\sigma}^k - (\nu\mu|^k \mathbf{P}_{\mu\nu}^k |\nu\sigma)^k / 2 \quad (\text{A.15b})$$

$$\mathbf{K}_{\mu\nu}^k \leftarrow \mathbf{K}_{\mu\nu}^k - (\mu\nu|^k \mathbf{P}_{\nu\sigma}^k |\sigma\nu)^k / 2 \quad (\text{A.15c})$$

$$\mathbf{K}_{\nu\nu}^k \leftarrow \mathbf{K}_{\nu\nu}^k - \left[(\nu\sigma|^k \mathbf{P}_{\sigma\mu}^k |\mu\nu)^k + (\nu\mu|^k \mathbf{P}_{\mu\sigma}^k |\sigma\nu)^k \right] / 2 \quad (\text{A.15d})$$

$$\mathbf{K}_{\mu\nu}^k \leftarrow \mathbf{K}_{\mu\nu}^k - (\mu\nu|^k \mathbf{P}_{\nu\nu}^k |\nu\nu)^k / 2 \quad (\text{A.16a})$$

$$\mathbf{K}_{\nu\nu}^k \leftarrow \mathbf{K}_{\nu\nu}^k - \left[(\nu\nu|^k \mathbf{P}_{\nu\mu}^k |\mu\nu)^k + (\nu\mu|^k \mathbf{P}_{\mu\nu}^k |\nu\nu)^k \right] / 2 \quad (\text{A.16b})$$

$$\mathbf{K}_{\mu\lambda}^k \leftarrow \mathbf{K}_{\mu\lambda}^k - (\mu\nu|^k \mathbf{P}_{\nu\nu}^k |\nu\lambda)^k / 2 \quad (\text{A.16c})$$

$$\mathbf{K}_{\lambda\nu}^k \leftarrow \mathbf{K}_{\lambda\nu}^k - (\lambda\nu|^k \mathbf{P}_{\nu\mu}^k |\mu\nu)^k / 2 \quad (\text{A.16d})$$

$$\mathbf{K}_{\mu\sigma}^k \leftarrow \mathbf{K}_{\mu\sigma}^k - (\mu\nu|^k \mathbf{P}_{\nu\mu}^k |\mu\sigma)^k / 2 \quad (\text{A.17a})$$

$$\mathbf{K}_{\nu\sigma}^k \leftarrow \mathbf{K}_{\nu\sigma}^k - (\nu\mu|^k \mathbf{P}_{\mu\mu}^k |\mu\sigma)^k / 2 \quad (\text{A.17b})$$

$$\mathbf{K}_{\mu\mu}^k \leftarrow \mathbf{K}_{\mu\mu}^k - [(\mu\nu|^k \mathbf{P}_{\nu\sigma}^k |\sigma\mu)^k \quad (\text{A.17c})$$

$$+ (\mu\sigma|^k \mathbf{P}_{\sigma\nu}^k |\nu\mu)^k] / 2 \quad (\text{A.17d})$$

$$\mathbf{K}_{\nu\mu}^k \leftarrow \mathbf{K}_{\nu\mu}^k - (\nu\mu|^k \mathbf{P}_{\mu\sigma}^k |\sigma\mu)^k / 2 \quad (\text{A.17e})$$